ATLAS NOTE
July 5, 2007

Data preparation and studies for the data streaming test

J.F. Arguin, A. Gaponenko, I. Hinchliffe, A. Holloway, M. Shapiro, L. Tompkins, S. Zenz

Abstract

This note describes the preparation and use of a special set of ATLAS Monte Carlo datasets, the 'streaming test' data. These datasets combine simulated events from many Standard Model processes, filtered by the ATLAS Level-1 and Level-2 trigger simulation and sorted into data streams. To provide a practical comparison of two alternative data streaming models, the same data are stored in both inclusive and exclusive streams. The reconstruction and analysis of these datasets is an important step towards readiness for data-taking, and a critical step in finalizing the choice of streaming model by the data streaming study group.

1 Introduction

Raw data streaming refers to the sorting of events based on their contents before they are permanently stored, in order to facilitate event reconstruction and user access. Typically, one of the first steps in data analysis involves selecting events that have satisfied certain trigger chains. Hence defining raw data streams based on triggers can reduce the load on ATLAS computing systems by reducing the number of events that each user must read.

It is natural and efficient to group triggers with similar signatures (such as electrons or missing energy) but different thresholds or multiplicities into the same stream. If data is sorted according to a given trigger menu, there are two obvious choices for the fate of events that pass multiple trigger chains assigned to different streams. These "overlapping" events can either be written to every stream for which they passed a trigger, or to one special stream containing all such events. These two scenarios are termed *inclusive* and *exclusive* streaming respectively, and each has advantages and disadvantages. The inclusive streaming model has the benefit that most users will only need one stream, and the possible loss of data in other streams will not affect their analyses. The exclusive streaming model has the obvious benefit that each event is stored only once[1)] at each stage of reconstruction.

The granularity of the data is an important aspect of the data streaming model. Naturally a single pp beam crossing is the smallest sensible unit of data, but it is not possible to monitor accelerator and detector conditions on an event-by-event basis. Instead all ATLAS file handling, event processing and analysis will treat a sequence of beam crossings spanning a time interval of order one minute as an indivisible unit called a *luminosity block*. The monitoring of detector conditions, the assignment of trigger prescales, and the transfer and reconstruction of events will only be performed for entire luminosity blocks. Moreover, all of the events from a given stream from any luminosity block will be available in a single reconstructed file. The luminosity block is the basis for calculating recorded luminosity, and hence when two streams are combined in an analysis, an identical set of luminosity blocks must be used from each stream.

To determine if the constraints of either streaming model are particularly inconvenient for users, a "streaming test" was devised. Ten short (30 minute) runs at an instantaneous luminosity near 10^{33} cm^{-2}s^{-1} were simulated by combining simulated events from release 11.0.X Monte Carlo RDO datasets. The simulated processes include W and Z boson production, single top quark and $t\bar{t}$ production, Drell-Yan pair production, and dijet production with jet p_T above 17 GeV. These events were stripped of all Monte Carlo truth information, mixed into a random order, filtered by the trigger simulation, and written to signature-based streams using both the inclusive and exclusive procedures. Each of the ten half-hour runs were divided into 30 luminosity blocks, which were reconstructed as indivisible units.

After reconstruction of the datasets, ATLAS collaborators were encouraged to analyze the data from relevant streams as if it were real data, and to report on the differences between the inclusive and exclusive models from their experiences. Since the technical difference between inclusive and exclusive streaming is mostly a question of bookkeeping, it was important that realistic prototypes of bookkeeping tools were available to users at this stage. For this reason, a luminosity block-based database of trigger prescales, delivered luminosity, detector livetime and stream status, and a TAG database for event filtering were built for the streaming test data.

In the following sections, we explain the construction and reconstruction of the streaming datasets, and the use of the bookkeeping and access tools to extract cross-sections from the streaming data.

---

1) Of course, datasets are generally replicated to different sites, so events are physically stored many times.

Stream	inclusive datasets	exclusive datasets
Jet	streamtest[_V1].00*.inclJet.digit.RDO.v1100399	streamtest.00*.exclJet.digit.RDO.v1100399
Electron	streamtest.00*.inclEle.digit.RDO.v1100399	streamtest.00*.exclEle.digit.RDO.v1100399
Muon	streamtest.00*.inclMuo.digit.RDO.v1100399	streamtest.00*.exclMuo.digit.RDO.v1100399
Photon	streamtest.00*.inclPho.digit.RDO.v1100399	streamtest.00*.exclPho.digit.RDO.v1100399
Tau/$\not\!\!E_T$	streamtest.00*.inclTau.digit.RDO.v1100399	streamtest.00*.exclTau.digit.RDO.v1100399
Overlap		streamtest.00*.Overlap.digit.RDO.v1100399

Table 1: Dataset naming scheme for inclusive and exclusive streams for the streaming test. There are ten datasets of each type (one per run). The inclusive jet stream dataset from run 0 is `streamtest_V1.004880.inclJet.digit.RDO.v11000399`; no other valid datasets are in the project streamtest_V1.

2 Preparing the raw streamed datasets

2.1 Structure of the streamed data

The entirety of the streaming test sample represents 18 pb^{-1} of data, collected in 10 runs of 30 minutes each. One reconstructed (ESD, AOD, TAG or ntuple) file from the streaming test contains all of the events from one or more luminosity blocks in one stream. In this section, we explain this structure, and describe how to access a particular subset of the streaming test data.

2.1.1 Streams

Five trigger categories are used to define the streams for the streaming test: electrons, muons, photons, tau leptons or missing energy, and jets. The detailed assignment of trigger signatures to streams is discussed in Section 2.2. In inclusive streaming mode, there are five datasets for each run, each containing all the events that passed triggers in one of the five categories. For the exclusive streaming, six datasets are written: one corresponding to each category above, and one for the overlap events. The resulting raw datasets are listed in Table 1.

2.1.2 Runs and luminosity blocks

One minute luminosity blocks are used for the streaming test. To encourage users of the datasets to test the bookkeeping tools, the average instantaneous luminosity for each luminosity block is varied. When the streaming test was conceived, it was assumed that ATLAS would begin new runs several times within an LHC fill, and that 30 minutes was a probable length for a single run. Therefore, 30 luminosity blocks make up each run in the streaming data.

Except for variations in delivered and recorded luminosity, the ten runs in the streaming test are functionally identical.

2.1.3 Structure of the raw datasets

When data-taking begins, events accepted by the ATLAS trigger will be written to disk by sub-farm output (SFO) modules (Figure 1). Several SFO's will operate in parallel, each writing a subset of the events from one luminosity block. The SFO's will perform the data streaming and hence will have N_{streams} files open at once, which will end up in different raw datasets. They will close files on luminosity block boundaries and transfer them to the Tier-0 for reconstruction. The files from all of the SFO's correspond-

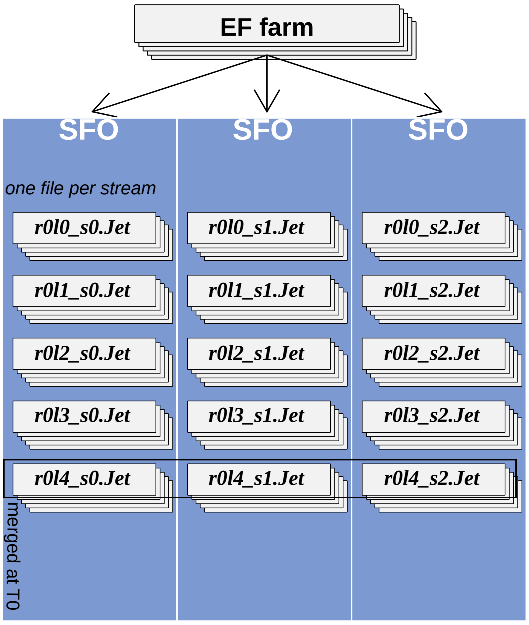

Figure 1: Illustration of raw data streaming at the SFO's. Each SFO writes to a new collection of files for every luminosity block. The files from the same stream and luminosity block are later merged.

ing to the same luminosity block will then be merged at the Tier-0, after which point luminosity blocks in each stream will be wholly contained within single files.

This aspect of data processing was simulated by running jobs in parallel to produce the streamed output. For every luminosity block, one batch job represented each SFO. These SFO-jobs performed the exclusive and inclusive streaming simultaneously, writing eleven files in total. Merging of files from the same stream and luminosity block was performed as part of the RDO $\rightarrow$ ESD reconstruction step.

The real SFO's receive and write data in bytestream format. The Monte Carlo data used for the streaming test was available in RDO (digit) format, and due to problems with the bytestream converters in release 12.0.3, the simulated SFO output is also in RDO format.

2.2 Event filtering and streaming

The trigger table used in data streaming determines the relative size of each stream, the fraction of events stored multiple times (by inclusive streaming) or in the overlap stream (by exclusive streaming), and ultimately the range of measurements and calibrations that can be performed with the collected data. The trigger table for the streaming test represents the first attempt to construct a trigger menu including prescales and stream assignments for high luminosity running (near 10^{33} cm^{-2}s^{-1}). Table 2 shows the trigger menu used for the streaming test, and the stream to which each trigger chain is assigned. It is roughly based on the trigger menu in earlier documents [**?**], although the prescale values in the table are partly motivated by technical constraints explained in Section 2.3.

2.2.1 Trigger simulation

The event mixing and trigger simulation used ATLAS software release 12.0.3. Only the Level-1 and Level-2 algorithms were run. The streaming test was the first dataset produced that used a full trigger menu, rather than specific chains or "slices." In most cases, however, the 12.0.3 Level-1+Level-2 default

chain name	**prescale**	**Level-1 condition(s)**	**Level-2 condition(s)**	Stream
jet25	2000000	JT15	L2_jet25	Jet
jet50	200000	JT15	L2_jet50	Jet
jet90	5000	JT50	L2_jet90	Jet
jet170	200	JT50	L2_jet170	Jet
jet300	20	JT100	L2_jet300	Jet
jet550	1	JT200	L2_jet550	Jet
4jet50	2	JT15(x4)	L2_jet50 (x4)	Jet
4jet110	1	JT50(x4)	L2_jet110 (x4)	Jet
sumet1000	1	L1_sumEt1T	pass-through	Jet
sumjet1000	1	L1_sumJt1T	pass-through	Jet
e15i	25	EM11	L2_e15i	Ele
e25i	1	EM18	L2_e25i	Ele
2e15i	1	EM11 (x2)	L2_e15i (x2)	Ele
e15i&mu10	1	MU10 AND EM11	L2_e15iL2_mu10	Ele
mu6	2	MU06	L2_mu6	Muo
mu20	2	MU20	L2_mu20	Muo
2mu10	1	MU10 (x2)	L2_mu10 (x2)	Muo
g20i	60	EM11	L2_g20	Pho
g60	1	EM50	L2_g60	Pho
2g20i	1	EM11 (x2)	L2_g20i (x2)	Pho
tau3w5i	10	HA30	L2_tau35i	Tau
tau35i&etmiss45	1	HA30 AND L1_Miss45	L2_tau35i	Tau
jet45&etmiss45	50	JT20 AND L1_Miss45	L2_jet45	Tau
jet70&etmiss70	1	JT50 AND L1_Miss70	L2jet70	Tau
etmiss200	1	L1_Miss200	pass-through	Tau
etmiss1000	1	L1_Miss1T	pass-through	Tau

Table 2: Level-1 and Level-2 trigger menu for the streaming test. Signature and threshold names are from the 12.0.3 mixing job and can be found in the XML files; the names assigned in STR-01 are occasionally different. The Event Filter was not simulated.

configuration of the trigger slices included trigger chains like those needed for the streaming test. Table 3 summarizes the changes we made to the 12.0.3 trigger configuration in order to implement the streaming trigger menu: most only required implementing new thresholds and can now can be obtained from more recent software releases. The XML files to configure this particular trigger menu in release 12.0.3 is in the CVS package offline/StreamingTest/StreamMix.

Trigger "slice"	**existing chains**	**chains added**	**other changes**
egamma	e15i, 2e15i, e25i, g20i, 2g10i, g60		combined e15 and mu10
muon	mu6	mu10, 2mu10, muo20, mu40	
jet	4j110	j25, j45, j70, j50, 4j50, j90, j170, j300, j550	combined jet triggers with met45 and met 70
tau	tau35i		combined with met45
missing/sum E_T		met45, met70, met200, met1000, sumet1000, sumjet1000	

Table 3: Changes to the 12.0.3 trigger configuration scripts.

To run all the trigger chains needed for the streaming test several developments to the trigger code were required, summarized in the following tags:

- Trigger/TrigConfiguration/TrigConfigSvc-00-00-08
- Trigger/TrigConfiguration/TrigConfHLTData-00-00-10
- Trigger/TrigConfiguration/TrigConfStorage-00-00-18
- Trigger/TrigRelease-00-03-85
- Trigger/TrigHypothesis/TrigJetHypo-00-01-10
- Trigger/TrigSteer/TrigResultBuilder-00-00-07-06
- Trigger/TrigSteer/TrigSteering-02-02-27-22

All other trigger software packages are the default packages in the 12.0.3 release. While the patches made it possible to run the event filtering and get reasonable results from each trigger slice, some features of the trigger simulation were not implemented or fully developed in this software release and could not be corrected. Specifically, these were problems with the muon endcap trigger, and the Level-2 $\not\!\!E_T$ and sum-E_T triggers (which were not implemented at all).

Because the output of the raw event mixing is in RDO format, for which no persistent representation of trigger information is available, most of the trigger information is lost when the streamed events are written – only the decision itself remains encoded in the event header. The trigger objects are important for data analysis, so this information is added back to the reconstructed streaming test data, by re-running the trigger simulation during reconstruction. However, reconstruction of the streaming data used a later software release than the raw event mixing, so the trigger simulation used a slightly different menu and configuration than the mixing and filtering jobs. The differences between the trigger menu for reconstructed data (menu STR-01, in trigger menu nomenclature) and the specialized streaming menu are summarized in Table 4. Both trigger decisions are accessible from the reconstructed files.

Trigger conditions	Changes to STR-01 with respect to 12.0.3
g20i	Absent in STR-01
e15i+mu10, tau35i+met45, jet(45,70)+met(45,70)	Combining triggers from different slices was not implemented in 12.0.6; these combinations must be constructed by hand from the relevant TriggerItems.
mu10,2mu10	Absent in STR-01. (mu10 was not in the streaming trigger menu, but is a component of e15imu10, so its absence is relevant for TAG and ntuple analysis.)

Table 4: Differences between the streaming trigger menu and STR-01 (the trigger menu configuration used to make ESD,AOD,ntuples and the TAGs.)

2.3 Event mixing

Rather than simulating 18 pb^{-1} of new data from many processes for the streaming test, we mixed existing simulated data. The mixed samples are based on a table of available simulated datasets which used Pythia 6.323, Herwig 6.507, AlpGen 2.05, and MC@NLO 3.10 as generators [1–4]. The datasets were all from the csc11 project and were simulated in release 11.0.4x. This simulation used a perfectly-aligned detector geometry, ATLAS-DC3-02.

These preexisting datasets fall into two categories, *sufficient* and *insufficient* luminosity, based on the cross section, generator filter efficiency, and events available for each Monte Carlo dataset. Insufficient datasets[2)] represented less than 18 pb^{-1} of a given process.

Because the streaming test sample is meant to reflect six hours of triggered ATLAS data, we had two choices for dealing with the insufficient datasets. When the shortfall was small, it was simplest to generate more data of the same type. For many of the low-p_T jet datasets, however, this was infeasible. Instead, we treated events from those samples as weighted by the ratio $\frac{18\ \text{pb}^{-1}}{\mathscr{L}_{\text{sample}}}$. These weights were *only* used to modify the trigger prescale values. For an event from Monte Carlo dataset Y with available luminosity $\mathscr{L}_Y$ pb^{-1}, the effective prescale (the inverse of the event's probability to be written out, given that it passed the trigger conditions) of a trigger chain with nominal prescale PS_{nom} is:

$$PS_{\text{eff}} = \begin{cases} 1 & \text{if } PS_{\text{nom}} \leq \frac{18}{\mathscr{L}_Y} \\ PS_{\text{nom}} \cdot \left(\frac{\mathscr{L}_Y}{18}\right) & \text{otherwise.} \end{cases} \tag{1}$$

In short, for insufficient samples, each applied prescale was as small as required to generate the appropriate number of passed events, or at least as small as possible. The resulting effective prescales were not required to be integers. As mentioned previously, many prescales used in the streaming test are larger than the likely values for a real 10^{33} $\text{cm}^{-2}\text{s}^{-1}$ trigger menu, in order to accommodate these corrections. We could not accommodate everything: we did not increase the lepton and photon trigger prescales in order to correct the shortfall of low p_T dijet events which occasionally pass them. Because of this, the QCD fake rate in the streaming test data is unrealistically low. Similarly, the prescales on high-threshold jet triggers could not always accommodate the tails of low-p_T dijet processes.

[2)] Here, we use the term dataset loosely, as we often combined datasets from different simulation tasks, and occasionally different generators, to augment the sample of events from a given process.

2.4 Event mixing input and output

Events were mixed using the MixingEventSelector in the Control/AthenaServices package. This service uses a list of basic EventSelectors, each initialized with a different input collection, and chooses among them according to a list of weights during event looping. Version AthenaServices-01-07-47 was used in the streaming test. In this version of the software, the MixingEventSelector ends the Athena job when any one of the EventSelectors runs out of input events. Because of this, fluctuations in the MixingEventSelector random sequence led to potentially large variations in the length of each mixing job. These were treated as variations in the detector livetime, as explained in Section 5.3.

Before writing the mixed events, two major changes were made to their contents. First, the Monte Carlo truth information was stripped by excluding Monte Carlo-related objects from the AthenaOutputStream. Second, the event header was rewritten to change the run number (which was previously the Monte Carlo dataset number and hence a "tag" of the event) and to include the proper run number, luminosity block, trigger and stream information. Replacing the old EventHeader in StoreGate required the following patches:

- Control/SGTools-00-01-01-04
- Control/StoreGate-02-15-20-07

Changing the persistent event header to permanently store the new trigger and stream information required the additional tag Event/EventAthenaPool-00-02-03.

The EventInfo object in the RDO thus contains the original Level-1 and Level-2 trigger decisions. These are stored as bit-masks in the TriggerInfo, which can be decoded using the trigger bit assignments documented in [5]. There are no provisions in the event header for the Level-1 trigger bits [6], but in order to provide this information we used the 32 bits of the previously unused field TriggerInfo::m_lvl1TriggerType as a bit-mask. Unfortunately, in release 12.0.6 this field is no longer unused, and the Level-1 trigger bits in the AOD have been overwritten.

3 Release 12.0.3 Trigger studies

As explained in Section 2, two versions of the trigger simulation code were run to prepare the streaming datasets. Here we describe studies of the release 12.0.3 trigger that we used to guide the preparation of the streamed raw data.

3.1 Trigger efficiency

We used Monte Carlo truth information to measure the trigger efficiency of the Level-1 and Level-2 triggers for single isolated leptons and jets, and for event missing E_T . For these studies, ntuples with truth and trigger information were created during the event mixing jobs, before events were filtered. Prescales were not used for these efficiency calculations. Hence, they represent unbiased measurements of the 12.0.3 trigger algorithm performance.

In Figures 2 and 3, we plot the efficiencies of some characteristic single-object triggers as a function of the truth p_T for electrons, photons, muons, and truth jets. For electrons, photons, and muons, stable particles with $p_T > 10$ and $|\eta| < 2.5$ were chosen. The sum of the energies of all other stable, interacting objects in a cone of 0.2 about the lepton or photon was required to be less than 15 GeV. For hadronically decaying tau leptons, we plot efficiency as a function of the true visible p_T(the tau p_T, minus the vector sum p_T of any neutrinos from its decay). We also plot the efficiency of the missing energy trigger as a function of the "truth $\not\!\!E_T$," which is the magnitude of the vector sum of all stable interacting truth particles within the bounds $|\eta| < 5.5$.

We verified that the efficiency plots have the desired threshold characteristics and represent reasonable efficiencies in all cases. The single muon trigger efficiency plateaus near 40% because of the problem with the Level-2 endcap trigger.

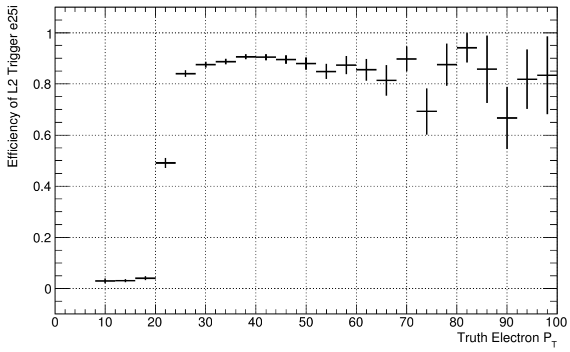

(a) Efficiency of `L2_e25i`

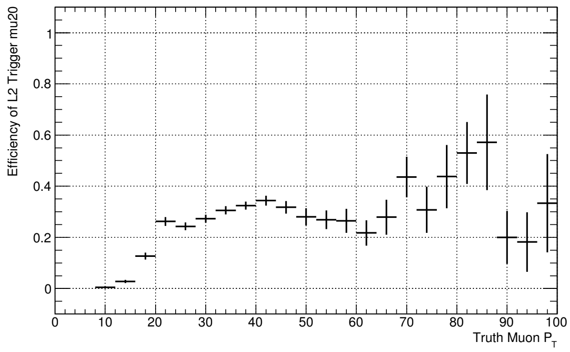

(b) Efficiency of `L2_mu20`

Figure 2: Efficiency of the e25i and mu20 triggers.

In Figure 4 we plot the efficiency for events to pass the tau35i trigger when the hadronic tau trigger candidate is matched to a truth electron (and not a tau lepton). This distribution, in light of the tau trigger efficiency plotted in 3(a), highlights a problem with the streaming trigger menu which is discussed in the following section.

3.2 Trigger overlaps

Real (truth-matched) electrons passing the electron trigger also frequently pass the tau and photon triggers in release 12.0.3, as indicated in Figure 4. Since e25i has a low prescale factor, the overlap between streams will be dominated by these coincidences[3)].

[3)] Although there is also a lot of similarity between jet and tau+$\not{E}_T$ trigger conditions, the low-p_T jet triggers are so heavily prescaled that the contribution of these tau/jet events to the overlap category will be relatively low.

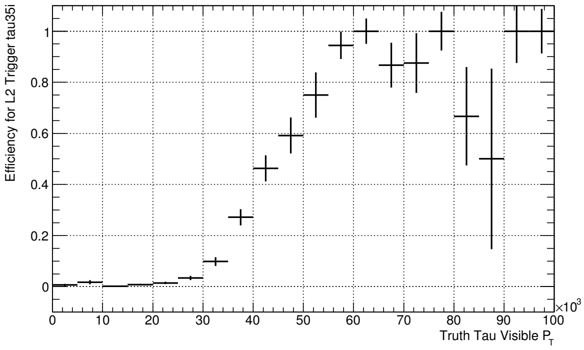

(a) Efficiency of `L2_tau35i`

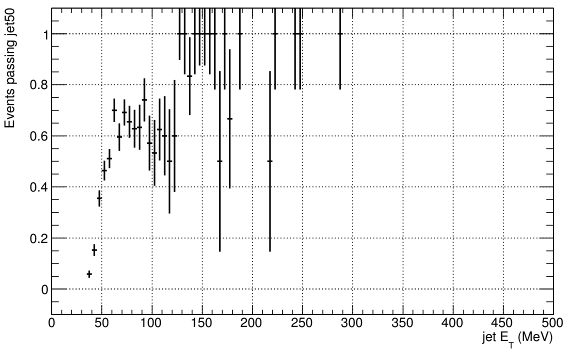

(b) Efficiency of `L2_jet50`

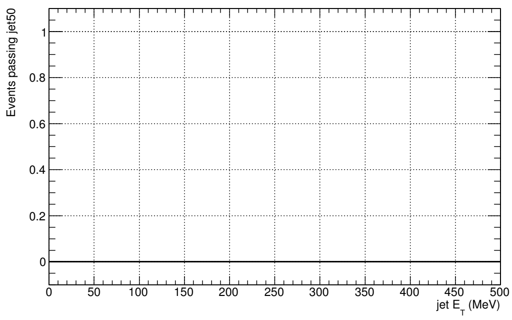

(c) Efficiency of `L2_je300`

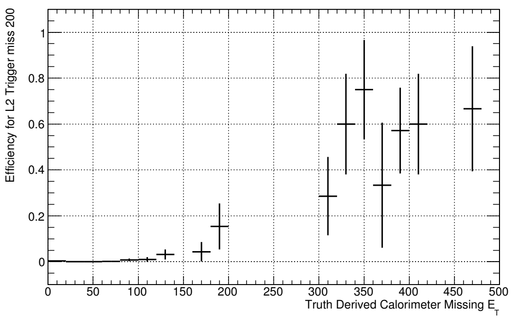

(d) Efficiency of `L2_x200`

Figure 3: Efficiency of the tau35i, $\not{E}_T > 200$, jet50‘, and jet300 triggers.

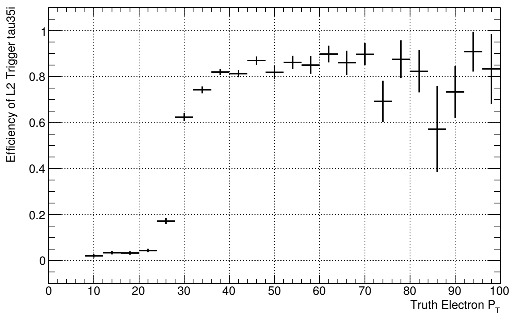

Figure 4: Efficiency of the tau35i trigger for electrons matched to the tau candidate, compared to the efficiency for hadronic taus.

Exclusive stream	fraction of streaming sample
Jet	22 %
Electron	35 %
Muon	20 %
Photon	5 %
Tau	6 %
Overlap	12 %

Table 5: Composition of the stream datasets resulting from the 12.0.3 trigger menu and out explicit trigger overrides.

In Figure 5 we plot the contribution of events from each stream to the overlap category, given the trigger-to-stream table presented in Section 2.2.1. Note that this figure over-represents the frequency of jet overlaps that would be expected if the streaming trigger table were applied to actual ATLAS data, because for the streaming test low p_T jet processes are weighted differently in different streams. However, the large contributions to the overlap stream from electron-tau and electron-photon overlaps are due to Monte Carlo processes with sufficient luminosity: these were weighted equally in all streams. It is undesirable in either streaming scenario for such a large fraction of real electron-containing events to fall into the overlap category.

Real ATLAS datasets will not necessarily have such large overlaps since the Event Filter will probably improve electron-photon and electron-tau separation. Moreover, the relative fraction of overlapping events will be less significant when the fake lepton rates are accurately represented, since fakes should not contribute to overlaps as frequently. In order to keep separate electron, photon, and tau streams in this streaming test, while preserving a reasonable overlap rate, we used truth information as a proxy for the Event Filter decision. For 90% of events in which both electron and photon or electron and tau triggers fired on a real electron, we explicitly vetoed the “wrong” trigger decision. This lowered the overall fraction of trigger overlaps from 20% to 12% of the entire sample, as indicated in Table 5.

4 Release 12.0.6 reconstruction of streaming data

The streaming test samples were reconstructed using the standard ATLAS production system. Raw datasets are registered using the output of the event mixing described in Section 2.4: these RDO datasets have version number `v12000399`, referring to a non-existent production cache. They are replicated from

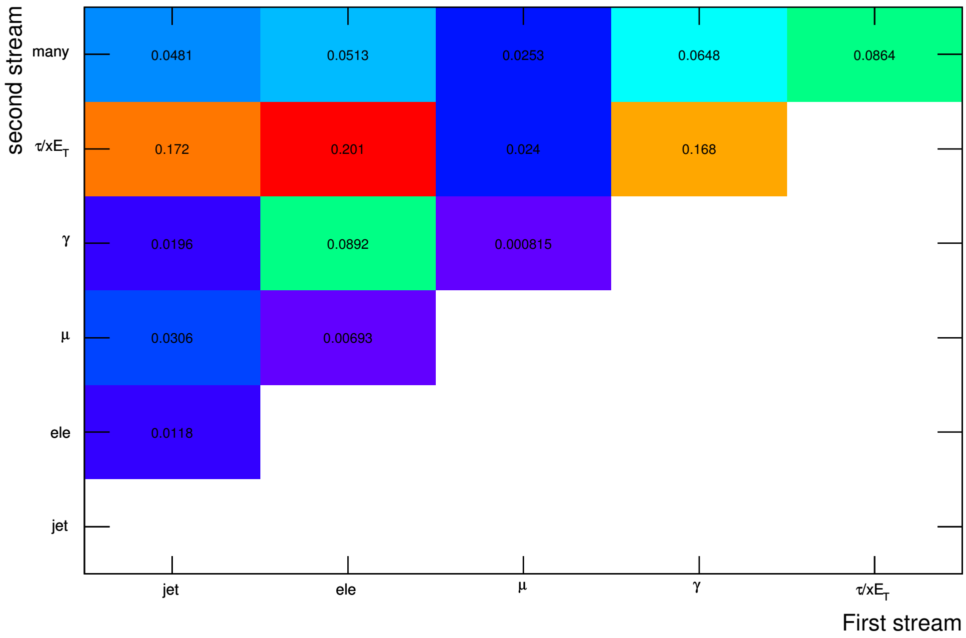

Figure 5: Summary of overlaps in the streaming table. The entries show the fraction of all overlap events stemming from each pair of categories.

BNL to CERN, LYON and NIKHEF. **(EDIT: Comments?)**

4.1 Reconstruction without Monte Carlo truth

In order for the streaming test reconstruction jobs to run without failure, several aspects of reconstruction had to be weaned of their dependence on Monte Carlo truth information [4]. Most cases were solved by setting an algorithm property to indicate that no truth information should be expected. The jobOption file `NoTruth.py` sets a global flag in reconstruction jobs for this purpose. Using the NoTruth file, production jobs in cache 12.0.6.4 and later can make ESD, AOD, SAN, HPTV, and TAG files from events without truth information.

4.2 "Empty" files

Because the number of processed events in each SFO-job fluctuated widely, there were occasional jobs that produced no output for a given stream. Of 33,000 anticipated output files, 61 were not created due to livetime fluctuations. This condition may occur during unstable detector conditions or for very small streams in real data, but no robust method exists for representing these cases of missing files in the current ATLAS computing model.

In addition, 89 files were lost during transfer between the batch system and permanent storage. When files are either lost in transfer or "legitimately" missing from a dataset, they cause the merge-and-reconstruction job to fail, and thereby remove the whole luminosity block from further processing

[4] This progress complements the development of the cosmics reconstruction chain as a step in preparation for 2008 data-taking, and so developers should continue to monitor reconstruction software performance on "realistic" data samples.

Stream	RDO (kB/event)	AOD (kB/event)	ESD (kB/event)
Jet	2319	173	1265
Ele	2198	76	723

Table 6: Average event sizes for the inclusive jet and electron data stream in the streaming test.

[5]. To repair this, the easiest solution was to replace the missing files with dummies. The dummy files could not be empty, since a POOL file must contain at least one event, or it will cause Athena to crash. We used dummy files containing a single minimum-bias event which passed no triggers and had no stream bits set, allowing users to avoid these extra events in their analysis.

4.3 Reconstructed file sizes

We emphasize that the relative sizes of different streams and the rate of streaming output when normalized to the 5 hours of data taking at 10^{33} cm^{-2}s^{-1} are not representative of our expectations for ATLAS data, because there are not enough low-p_T QCD events passing other triggers. Individual raw data files are uncharacteristically small for the same reasons. Nevertheless, because of the mixing and trigger filtering, the streaming test data is probably a better indicator of average event sizes for real data than the $t\bar{t}$ simulation samples that have been used previously. Table 6 describes the average size of events in the streaming data sample, estimated from one luminosity block in the ESD, ten in the AOD, and 20 in the RDO.

5 Using metadata and database tools for analysis

The construction of online databases for trigger configuration and run conditions and luminosity information, and of the offline TAG database, are related to the data streaming test and should be documented elsewhere [7, 8]. Prototypes of each of these databases and of the luminosity-block metadata storage format were prepared for the streaming test. In this note we simply summarize how these prototype tools can be used to analyze the streaming datasets.

5.1 Luminosity-block metadata

Since the luminosity block is the basis of sample luminosity calculations, every analysis must be able to access the complete list of luminosity blocks which were processed in deriving the final analysis samples (even if no events from those blocks remain in the final samples). Moreover, as information on detector conditions evolves and the database of "bad" luminosity blocks is updated, physicists must be able to remove the events from those bad blocks from their event collections *and* their list of used blocks at the same time. The metadata container which can be added to POOL files helps accomplish this, and tools in the LumiBlockComps package maintains the luminosity metadata as files are processed by Athena [9].

Ideally this file-resident luminosity metadata is created at the first stage of reconstruction, when raw files are merged to form complete luminosity blocks. It should be updated and written into the output files at each stage of reconstruction when input files are skimmed or merged. However, the metadata collection was not implemented when the streaming test event processing began, and no persistency format for the luminosity block metadata exists in release 11 or 12. For this reason, users of streaming

[5] This may in fact be the desired behavior when the file is lost due to corruption or failed transfer, because it incidentally removes incomplete luminosity blocks from reconstructed datasets. However, for the streaming test we correcd the absence of files from both causes, so corrupted luminosity blocks must be removed in a later step.

test data must use release 13 tools to extract the metadata from an (unfiltered) collection of events from their analysis, and write it into a file. Assuming they intend to analyze events in the better-validated release 12 format, their metadata file must be external to the streaming test AOD or ntuple files.

An example script for this purpose, `MakeMetaDataNT_jobOptions.py`, can be found in the LumiBlock/LumiBlockComps package. It currently runs in release 13.X.0 nightlies.

5.2 Selecting with TAGs

The TAG database is available online and as a collection of ROOT files in DQ2 [7]. This tool can be used to quickly summarize the contents of a database, or to select events for analysis. Queries can involve run or luminosity block number, as well as trigger and analysis cuts: care must be taken when mixing the two kinds of selections, since the former cuts must also be applied to the metadata collecion.

The first version of the TAG database only tabulated the trigger information from the 12.0.6 STR-01 trigger (*i.e.*, the decision objects in the reconstructed files). Queries on these trigger decisions will occasionally reveal discrepancies between the stream containing an event and the triggers passed by an event, for two reasons: first, none of the decisions in the STR-01 table are prescaled, and second, some 12.0.6 trigger algorithms are slightly different from their 12.0.3 counterparts[6)] Moreover, because not every trigger decision in the streaming table is represented by a *signature* in STR-01, it is impossible to check some of these with a TAG query. In these cases, the trigger objects in the AOD can generally be used.

5.3 Querying the luminosity

A luminosity query requires the luminosity database [8] and the luminosity block metadata for an event collection. For the data streaming test, a luminosity database was constructed from the log files of the event mixing stage, which reported the nominal instantaneous luminosity at the beginning of the job, and the total number of events processed at the end of the job. The log file also contained the number of events passing each trigger signature in the job, and reported the prescale settings for all of the trigger chains.

For the first database prototype, an ad-hoc live-time fraction was used for all luminosity blocks. In an updated prototype, the live time is derived as follows: when, due to an input shortfall, a job did not process the desired number of events, the fraction of events that were not processed is used as a "deadtime" correction.

Data transfer failures are not available in the Athena log files and not recorded in the RunLumiDB prototype. For real data-taking, the failed transfer (or corruption) of one or more SFO's output files must be reported, so that the remaining luminosity block fragments are not used for analysis. Note that a single data transfer failure will only corrupt the luminosity blocks in one stream, unless the failure affects an overlap stream file.

A luminosity calculation tool without the corrupted luminosity block correction is available in release 13 nightly builds. The tool, `LumiCalc.py`, is found in the LumiBlock/LumiBlockComps package and interfaces to the database prototype and the metadata store in POOL or ROOT files. It can be used with the standalone metadata files produced using the job-options script above:

```
[you@lxplus] source ~/cmthome/setup.csh -tag=13.X.0,rel_4,setup
[you@lxplus] LumiCalc.py --trigger=L2_met200 myMetaDataColl.root
<some warnings ... >
>== Trigger  : L2_e25i
IntL (nb^-1) :       488.88
```

[6)] In the case of single jet triggers, it appears that the 12.0.6 configuration is very different than the 12.0.3 configuration; thresholds from the STR-01 decision can be as much as 50% lower than nominal.

```
L1/2/3 accept:         9787     4281        0
Livetime     :     582.0000
Good LBs     :           10
BadStatus LBs:            0
```

6 Conclusions

Basically, this part will be the to-do list:

- Interaction between tag queries and luminosity metadata
- Home for file transfer deadtime information
- flexible reconstruction system missing files

Fixes from user feedback:

- fix jets, jet triggers
- FDR: more event generation for more productive stream balance
- FDR: merge similar objects like electron and photon

(EDIT: other issues?)

References

[1] Sjostrand, Torbjorn and Lonnblad, Leif and Mrenna, Stephen and Skands, Peter, PYTHIA 6.3: Physics and manual, 2003.

[2] G. Corcella and I. G. Knowles and G. Marchesini and S. Moretti and K. Odagiri and P. Richardson and M. H. Seymour and B. R. Webber, JHEP **0101** (2001) 010.

[3] M. L. Mangano and M. Moretti and F. Piccinini and R. Pittau and A. D. Polosa, JHEP **0307** (2003) 001.

[4] S. Frixione and B. R. Webber, The MC@NLO Event Generator, 2002.

[5] ”StreamingTestForUsers, https://twiki.cern.ch/twiki/bin/view/Atlas/StreamingTestForUsers”.

[6] C. Bee, D. Francis, L. Mapelli, R. McLaren, G. Mornacchi, J. Petersen, and F. Wickens, The raw event format in the ATLAS Trigger and DAQ, ATL-D-ES-0019v31.

[7] ”TAG, https://twiki.cern.ch/twiki/bin/view/Atlas/TAG”.

[8] ”StreamingTestForUsers, https://twiki.cern.ch/twiki/bin/view/Atlas/”.

[9] Costanzo, D and others, Metadata for ATLAS, Apr 2007.